

# Defunctionalizing Push Arrays

Josef Svenningsson    Bo Joel Svensson

Chalmers University of Technology

[josefs, joels]@chalmers.se

## Abstract

Recent work on domain specific languages (DSLs) for high performance array programming has given rise to a number of array representations. In Feldspar and Obsidian there are two different kinds of arrays, called Pull and Push arrays and in Repa there is an even higher number of different array types. The reason for having multiple array types is to obtain code that performs better. Pull- and Push arrays, that are present in Feldspar and Obsidian, provide this by guaranteeing that operations fuse automatically. It is also the case that some operations are easily implemented and perform well on Pull arrays, while for some operations, Push arrays provide better implementations. Repa has Pull arrays, called *delayed* arrays for the same reason and so-called *cursorred* arrays which are important for performance in stencil operations. But do we really need to have more than one array representation? In this paper we derive a new array representation from Push arrays that have all the good qualities of Pull- and Push arrays combined. This new array representation is obtained via defunctionalization of a Push array API.

**Categories and Subject Descriptors** CR-number [subcategory]: third-level

**General Terms** term1, term2

**Keywords** keyword1, keyword2

## 1. Introduction

Recent developments in high performance functional array programming has given rise to two complementary array representations, Pull- and Push arrays. Pull is the traditional type `Idx -> Value`; a function from index to a value. Push arrays, on the other hand, are programs parameterised on a write function that take an index and a value. Pull- and Push array implementation details are shown in sections 2.2 and 2.4. These two representations have many advantages:

- They are easily parallelizable. It is straight forward to generate efficient, parallel code from Pull- and Push arrays, making them suitable for inclusion in high performance DSLs.
- They allow for a compositional programming style. It is easy to formulate high level, reusable combinators operating on Pull- and Push arrays.

- They support fusion. When composing two array functions the intermediate array is guaranteed to be fused away and not allocated in memory at runtime. Having such guarantees makes the compositional programming style particularly attractive, as it comes without any performance overhead.

The reason there are two types of arrays is that they complement each other. Some operations, like indexing, can be implemented efficiently for Pull arrays but not for Push arrays. Other operations, such as concatenation, are more efficient on Push arrays compared to Pull arrays.

But why do we need two types of arrays? It would certainly be easier for the programmer if there was only a single type to keep track of.

This paper presents a unified array library which inherits the benefits of Pull- and Push arrays, yet has only a single type of array.

- We present a single array type which subsumes both Pull- and Push arrays (section 5).
- We show how to derive our new array type by applying *defunctionalization* on push arrays (section 4).
- Our array library can support all known safe operations on Pull- and Push arrays. See section 7.2 for a discussion.
- We present our new array library in the context of an embedded DSL (section 2.1) because compiled DSLs easily provide fusion guarantees. Section 7.1 outlines how to achieve fusion for our library in the context of Haskell.

Before presenting the contributions we will give an introduction to Pull- and Push arrays in section 2.2 and 2.4. We will also review defunctionalization in section 3.

## 2. Background

In this section we will present enough background material to make the paper self-contained. Nothing here is new material.

### 2.1 Preliminaries

In this paper we present Pull- and Push arrays as part of a small code generating embedded language, a compiled EDSL [15]. The embedded language is compiled into a small imperative language with for loops, memory operations (allocate, write) and conditionals, a simple C like language. The data type `Code` below is used to represent compiled programs.

```
type Id = String

data Code = Skip
          | Code :>>: Code
          | For Id Exp Code
          | Allocate Id Length
          | Write Id Exp Exp
          | Cond Exp Code Code
```

[Copyright notice will appear here once 'preprint' option is removed.]

```

-- Memory operations
write :: Expable a => CMMem a -> Expr Int -> a -> CM ()

cmIndex :: Expable a => CMMem a -> Expr Int -> a

-- for loop
for_ :: Expable a => Expr Int -> (a -> CM ()) -> CM ()

-- conditionals
cond :: Expr Bool -> CM () -> CM () -> CM ()

```

**Figure 1.** Memory operations and flow control

There are also scalar expressions in the target language. These are represented by a data type called `Exp`. Here all the usual arithmetic operations, indexing and expression level conditionals are found. Functions in the embedded language that do not mesh well with the existing type class system are given a name ending with an underscore, such as `mod_`, and conditionals are expressed using an `ifthenelse` function.

```

data Value = IntVal Int
           | FloatVal Float
           | BoolVal Bool

data Exp = Var Id
         | Literal Value
         | Index String Exp
         | Exp :+: Exp
         | ...
         | LEq Exp Exp
         | Min Exp Exp
         | IfThenElse Exp Exp Exp

```

Phantom types provide a typed interface to expressions, while it allows a simple (not GADT based) implementation of expressions. Again this follows the presentation of the seminal “Compiling Embedded Languages” [15].

```
data Expr a = E {unE :: Exp}
```

An `Expable` class sets up conversions to and from the `Exp` data type. Compilation will require that data elements are an instance of this class.

```
class Expable a where
  toExp :: a -> Exp
  fromExp :: Exp -> a

```

Compilation down to this language is performed within a monad, a *Compile Monad* (`CM`).

```
newtype CM a =
  CM (StateT Integer (Writer Code) a)
  deriving (Monad,
           MonadState Integer,
           MonadWriter Code)

```

`CM` is a `Code` generating monad implemented as a writer/state combination. The state is used to construct new identifiers, while the writer accumulates generated code. Figure 1, shows type signatures of some essential functions that are used in the rest of the paper. The `CMMem` data type represents arrays in memory and is defined as a length and an identifier.

The `Code` data type can be seen as a deeply embedded core language. For the purposes of this paper, we consider `Code` the compiler result even though more steps are needed to actually run the code on an actual machine.

```

map :: (a -> b) -> Pull a -> Pull b
map f (Pull l ixf) = Pull l (f . ixf)

index :: Pull a -> Ix -> a
index (Pull _ ixf) i = ixf i

zipWith :: (a -> b -> c) -> Pull a -> Pull b -> Pull c
zipWith f (Pull l1 ixf1) (Pull l2 ixf2) =
  Pull (min l1 l2) (\i -> f (ixf1 i) (ixf2 i))

halve :: Pull a -> (Pull a, Pull a)
halve (Pull l ixf) = (Pull l2 ixf, Pull (l - l2) ixf')
  where l2 = l `div` 2
        ixf' i = ixf (i `:+` l2)

```

**Figure 2.** Examples of operations on Pull arrays

On top of the deeply embedded `Code` language, the two array representations (`Pull` and `Push`) are implemented as shallow embeddings [23]. Lengths and indices for these arrays are represented by the following types.

```
type Length = Expr Int
type Ix = Expr Int

```

The following sections will show the implementation- and properties of `Pull` and `Push` arrays.

## 2.2 Pull Arrays

A well known, and often used, representation of arrays is as a function from index to value. A value at a given index of an array is computed by applying the function to that given index. That is, pulling at an index provides a value. We call arrays that are implemented using this representation *Pull* arrays.

```
data Pull a = Pull Length (Ix -> a)
```

Examples of operations on `Pull` arrays are shown in Figure 2. A notable thing about `Pull` arrays and its functions is that they are non-recursive. This has the effect that it is very easy to fuse `Pull` arrays such that any intermediate array is removed, in particular in the context of embedded DSLs [23]. As an example, consider the expression `zipWith f (map g a1) (map h a2)` where `a1` and `a2` are `Pull` arrays. When this expression is evaluated by the Haskell runtime the definition of the individual functions will be unfolded as follows:

```

zipWith f (map g (Pull l1 ixf1)) (map h (Pull l2 ixf2))
=> zipWith f (map g (Pull l1 ixf1)) (Pull l2 (h . ixf2))
=> zipWith f (Pull l1 (g . ixf1)) (Pull l2 (h . ixf2))
=> Pull (min l1 l2) (\i -> f (g (ixf1 i)) (h (ixf2 i)))

```

The end result is a single `Pull` array and all the intermediate arrays have been eliminated.

`Pull` arrays are not named areas of memory containing data. Instead, a `Pull` array describes how the values can be computed. However, sometimes it is necessary to be able to compute and store the array as data in memory, for example to facilitate sharing of computed results. This can be done by adding a `force` primitive to the `Pull` array API.

```
force :: Expable a => Pull a -> CM (Pull a)
```

The `force` function is monadic (in the `CM` monad). Using the `force` function accumulates code into the writer part of the monad. The code accumulated is a program that iterates over, computes and stores all the elements of the input array to memory. Using `force`, is the only way to stop operations on `Pull` arrays from fusing. The

array returned from `force` contains the exact same data as the input, only they are now represented in memory.

### 2.3 Push Arrays

Push arrays are a complement to Pull arrays, first introduced into Obsidian [10]. Since then, Push arrays have also been implemented in Feldspar<sup>1</sup>, in Nikola<sup>2</sup> and in meta-repa [5]. Push arrays are also used in reference [20], as part of an embedded language for stream processing.

Push arrays were introduced in Obsidian and Feldspar in order to deal with very specific performance issues. In particular, array concatenation and interleaving introduces conditionals in the Pull array indexing function. When forcing an array, such conditionals can lead to bad performance on both CPUs and GPUs.

The example below shows a situation that may occur when working with pull arrays and concatenation. The code on the left executes a conditional in every iteration of the loop body. To the right, the loop is split into two separate loops, neither containing a conditional.

<pre>for i in 0..(m + n-1)   data[i] = if (i &lt; m)             then ...             else ...</pre>	<pre>for i in 0..(m-1)   data[i] = ... for i in 0..(n-1)   data[i+m] = ...</pre>
--	--

Another example, that occurs when flattening an array of pairs, is a loop that executes twice as many times as the array of pairs is long. In each iteration it selects the first or second component of the pair depending on whether the index is even or not.

<pre>for i in 0..(2*n-1)   data[i] = if even(i)             then fst(...)             else snd(...)</pre>	<pre>for i in 0..(n-1)   data[2*i] = fst(...)   data[2*i+1] = snd(...)</pre>
---	--

When working with Pull arrays, the loop structures to the left in the examples above are obtained. However, the code on the right is preferred. By switching to Push arrays the loop structures on the right can also be implemented.

Push arrays move the responsibility of setting up the iteration schema from the consumer (as with Pull arrays) to the producer. This provided, concatenation, interleaving and pair flattening can be given more efficient Push array implementations.

### 2.4 A Push Array Library

In this section, Push arrays are added to the embedded language. Just like the Pull arrays, Push arrays are added as a shallow embedding.

```
data Push a =
  Push ((Ix -> a -> CM ()) -> CM ()) Length
```

The Push array is a higher order function whose result is a monadic computation. As input, this higher order function takes a *write-function* (`Ix -> a -> CM ()`), that represents a way to write an element to memory. The Push array can then use this write-function any number of times. There is also a connection between this representation of Push arrays and continuations (see figure 3).

Figure 4, lists the Push array API used as basis for the defunctionalization in the upcoming sections. The selection of functions in the API is based on our experience with Push arrays from working with embedded languages.

Note that there is no arbitrary permutation, `ixMap`, in the library.

<sup>1</sup> [github.com/Feldspar/feldspar-language](https://github.com/Feldspar/feldspar-language)

<sup>2</sup> [github.com/mainland/nikola/blob/master/src/Data/Array/Nikola/Repr/Push.hs](https://github.com/mainland/nikola/blob/master/src/Data/Array/Nikola/Repr/Push.hs)

Perceptive readers will see that this is related to continuations as follows:

```
data Cont r a = (a -> r) -> r
```

Then Push arrays can be implemented as:

```
type Push m a = Cont (CM ()) (Ix,a)
```

In this way we can understand Push arrays as computations which non-deterministically generates index-value-pairs, implemented using a continuation monad.

Figure 3.

```
ixMap :: (Ix -> Ix) -> Push a -> Push a
ixMap f (Push p l) =
  Push (\k -> p (\i a -> k (f i) a)) l
```

This, somewhat naive, `ixMap` function is dangerous and can lead to uninitialized elements in the resulting array. Instead we have chosen to argue for using a set of fixed permutations, such as `reverse` and `rotate`. For the discussion in section 5, it is also important that these permutations are invertible.

### 2.5 Pull and Push Array Interplay

Pull and Push arrays complement one another and when programming it is nice to have both. Some functions are efficient and intuitive on Pull arrays. Function such as `zipWith` and `pair`, are *Pully* in nature, while `unPair` is more efficient on Push arrays, call it *Pushy*.

There is also a connection between the concepts of *scatter/gather* operations and Push/Pull arrays. Pull arrays allow implementation of operations that gather, while Push arrays enable scatter operations.

Converting a Pull array to a Push array is cheap and is also subject to fusion. The function `push` below implements this conversion.

```
push (Pull n ixf) =
  Push (\k -> for_ n $ \i ->
        k i (ixf i)) n
```

However, converting a Push array to a Pull array requires computing and storing all elements to memory. The function `pull`, implemented below, is an example of this.

```
pull :: Push a -> CM (Pull a)
pull (Push p n) =
  do arr <- allocate n
     p $ write arr
     return $ Pull n (\i -> cmIndex arr i)
```

This encourages the following pattern when programming with Pull- and Push arrays: A function takes one or several Pull arrays as arguments. These arrays are split apart and some processing is done on the individual parts. As a final step the arrays are assembled together again and produces a Push array. This push array can then be stored to memory which then can be read back as a Pull array again, if needed. Since memory accesses are an important factor in application performance on many platforms, the number of Push to Pull conversions can be used as a crude indicator of performance. Few such conversions is likely to be better.

An example of this pattern can be seen below in the function `halfCleaner` below. For an array of size 8 it performs compare and swap on pairs of elements at the following positions (0,4), (1,5), (2,6) and (3,7). This is achieved by first splitting the array in half. Then the two halves are zipped together and compare and swap is performed on the pairs. Finally, the new halves are unzipped and concatenated together, thereby producing the final array.

```

-- Array creation

generate :: Expable a
  => Length -> (Ix -> a) -> Push a
generate n ixf = Push (\k -> for_ n $ \i ->
                      k i (ixf i))

-- Map

map :: Expable a
  => (a -> b) -> Push a -> Push b
map f (Push p l) =
  Push (\k -> p (\i a -> k i (f a))) l

imap :: Expable a
  => (Ix -> a -> b) -> Push a -> Push b
imap f (Push p l) =
  Push (\k -> p (\i a -> k i (f i a))) l

-- Permutations

reverse :: Push a -> Push a
reverse (Push p n) =
  Push (\k -> p (\i a -> k (n - 1 - i) a)) n

rotate :: Length -> Push a -> Push a
rotate d (Push p n) =
  Push (\k -> p (\i a -> k ((i + d) `mod` n) a)) n

-- Combining Push arrays

(++) :: Expable a => Push a -> Push a -> Push a
(Push p1 l1) ++ (Push p2 l2) = Push r (l1 + l2)
  where r k = do p1 k
                p2 (\i a -> k (l1 + i) a)

interleave :: Push a -> Push a -> Push a
interleave (Push p m) (Push q n) = Push r l
  where r k = do p (\i a -> k (2*i) a)
                q (\i a -> k (2*i+1) a)
        l = (2 * (min_ m n))

-- Create Push array from data in memory

use :: Expable a => CMMem a -> Length -> Push a
use mem l = Push p l
  where
    p k = for_ l $ \ix ->
          k ix (cmIndex mem ix)

toVector :: Expable a => Push a -> CM (CMMem a)
toVector (Push p l) =
  do
    arr <- allocate l
    p $ write arr
    return arr

```

**Figure 4.** Our Push array API. The functions are representative of what we use when programming with Push arrays in Obsidian and Feldspar.

A half cleaner is an integral part in bitonic sorts and a similar pattern can be used to implement the butterfly network in FFT.

```
swap (a,b) = IfThenElse (a <: b) (a,b) (b,a)
```

```
halfCleaner :: Pull (Expr a) ->
             Push (Expr a)
halfCleaner =
  uncurry (++) . unzip . map swap . uncurry zip . halve
```

Note that the splitting and zipping must be done on a Pull array as those operations cannot be efficiently implemented using Push arrays. Mapping and unzipping can be done on either representation but in this case it is done on Pull arrays. Only the last step, concatenation, results in a Push array.

While Pull- and Push arrays work well together and form powerful abstractions for expression and control of computations, having just one array representation is alluring.

### 3. Defunctionalization

Defunctionalization is a program transformation introduced by Reynolds [22]. It is used to convert functions to first order data types which means it can be used as a way to implement higher order languages. We work in a typed setting and will follow the presentation of Pottier and Gauthier [21].

To illustrate defunctionalization we present a small example, originally due to Olivier Danvy [13]. The following program flattens trees into lists. A naïve flattening function has a worst case quadratic time complexity, because of nested calls to append. The version below is linear by using the standard trick of John Hughes [18] to represent lists as functions from lists to lists.

```

data Tree a = Leaf a
            | Node (Tree a) (Tree a)

cons :: a -> ([a] -> [a])
cons x = \xs -> x : xs

o :: (b -> c) -> (a -> b) -> a -> c
f 'o' g = \x -> f (g x)

flatten :: Tree t -> [t]
flatten t = walk t []

walk :: Tree t -> ([t] -> [t])
walk (Leaf x) = cons x
walk (Node t1 t2) = walk t1 'o' walk t2

```

The function walk is currently higher order. One way to make it a first order function is to eta-expand it and inline cons and o. We will instead use defunctionalization to make a first order version of the whole program.

Defunctionalization works in three steps. First, the function space we wish to defunctionalize is replaced by an algebraic data type. Second, lambda abstractions are replaced by constructors in that data type. And third, function application is replaced by a new function which interprets the algebraic data type such that the semantics of the program is preserved.

For the program above we create a new data type Lam a which replaces the functions from lists to lists, i.e. [a] -> [a]. There are two lambda abstractions which will be turned into constructors of the Lam data type. They are underlined in the code above. When replacing lambda abstractions by constructors it is important to capture the free variables as arguments to the constructor. In the lambda abstraction occurring in cons there is one free variable, x. We will therefore create a constructor Lam, which takes one argument of type a. Similarly for the abstraction in o, there are two free variables f and g. Since they are function arguments,

they will turn into elements of the data type `Lam a` and hence the constructor to replace the lambda abstraction will have two recursive arguments.

Finally, we need to create the function `apply` which interprets the constructors. Since our data type `Lam a` represents functions over lists the `apply` function will have type `Lam a -> ([a] -> [a])`. The function `apply` is defined with one case per constructor, where the result is the corresponding lambda abstraction in the original program which the constructor replaced. Having defined `apply`, we also need to insert it at the appropriate places in the program. In our case it will be in the function `flatten` which applies the function from `walk`. We will also need it in the definition of `apply` when composing the two defunctionalized functions.

The final result of defunctionalizing our example program can be seen below.

```
data Lam a = LamCons a
           | Lam0 (Lam a) (Lam a)

apply :: Lam a -> [a] -> [a]
apply (LamCons x) xs = x : xs
apply (Lam0 f1 f2) xs = apply f1 (apply f2 xs)

cons_def :: a -> Lam a
cons_def x = LamCons x

o_def :: Lam a -> Lam a -> Lam a
o_def f1 f2 = Lam0 f1 f2

flatten_def :: Tree t -> [t]
flatten_def t = apply (walk_def t) []

walk_def :: Tree t -> Lam t
walk_def (Leaf x) = cons_def x
walk_def (Node t1 t2) = o_def (walk_def t1) (walk_def t2)
```

A key observation in this example is that we have performed *local* defunctionalization. If the above code was contained in a module which only exported the `Tree` data type and the `flatten` function, then the defunctionalization we performed could be done completely independently of the rest of the program. So even though defunctionalization is in general a whole program transformation, there are programs where it can be applied locally. We will make use of this fact in the remainder of the paper.

## 4. Defunctionalizing Push Arrays

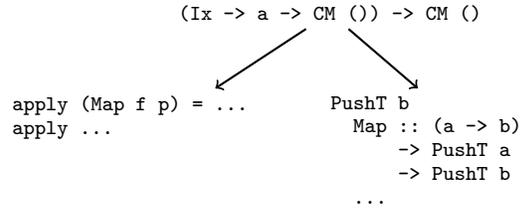
We now turn to applying defunctionalization in Push arrays. There are several potential functions to defunctionalize in the definition of Push arrays. We are going to focus on the outermost function, which takes the write function as argument and returns `CM ()` as result, as showed in figure 5. Applying defunctionalization to this function result in a data type (`PushT`) with constructors that represents our operations on Push arrays. We also obtain an `apply` function, an interpreter for our new Push array language.

We show how to defunctionalize `map` and `(++)` in full. The other functions from the API follow the same procedure. The procedure begins by investigating the body of the `map` function.

```
map :: (a -> b) -> Push a -> Push b
map f (Push p l) =
  Push (\k -> p (\i a -> k i (f a))) l
```

The free variables in the underlined body is `p` and `f`. These two will be arguments to our `Map` constructor that is added to the `PushT` data type.

```
data PushT b where
  Map :: (a -> b) -> PushT a -> PushT b
```



**Figure 5.** Using defunctionalization, we go from a higher order function to a data type and apply function.

The implementation of the `map` function exposed to the programmer is simply the `Map` constructor.

```
map :: (a -> b) -> PushT a -> PushT b
map = Map
```

The `apply` function is given from the body of the original `map` function, `\k -> p (\i a -> k i (f a))`. However, `p` is no longer a function. So an application of `apply` is needed to make the types match up.

```
apply :: (...) => PushT a -> (Ix -> a -> CM ()) -> CM ()
apply (Map f p) = \k -> apply p (\i a -> k i (f a))
```

Note that in our original definition of Push arrays, the length was stored with the array. The defunctionalized definition of Push arrays does not have this associated length. Instead the length is found by traversing the `PushT` data type. This is just a stylistic choice, to make the presentation cleaner.

To defunctionalize `(++)` we begin as with `map` by looking at the body of the function. In this case, essentially the `where` clause.

```
(++) :: Monad m => Push a -> Push a -> Push a
(Push p1 l1) ++ (Push p2 l2) = Push r (l1 + l2)
  where r k = do p1 k
            p2 (\i a -> k (l1 + i) a)
```

The free variables are `p1`, `p2` and `l1`. The constructor `Append` is chosen to represent this operation and is added to the `PushT` data type.

```
data PushT b where
  Map :: (a -> b)
      -> PushT a
      -> PushT b
  Append :: Ix
          -> PushT b
          -> PushT b
          -> PushT b
```

The `apply` functions gets a new case for `Append`.

```
apply :: (...) => PushT a -> (Ix -> a -> m ()) -> m ()
apply (Map f p) =
  \k -> apply p (\i a -> k i (f a))
apply (Append l p1 p2) =
  \k -> apply p1 k >>
        apply p2 (\i a -> k (l + i) a)
```

As with `map` the new `(++)` function is implemented directly from the `Append` constructor.

```
(++) :: PushT a -> PushT a -> PushT a
p1 ++ p2 = Append (len p1) p1 p2
```

This procedure is repeated for all operations in our API resulting in the data type and apply function shown in figure 6.

<pre> data PushT b where   Generate :: Expable b             =&gt; Length             -&gt; (Ix -&gt; b)             -&gt; PushT b    Use :: Expable b        =&gt; Length        -&gt; CMMem b        -&gt; PushT b    Map :: Expable a        =&gt; (a -&gt; b)        -&gt; PushT a        -&gt; PushT b    IMap :: Expable a         =&gt; (Ix -&gt; a -&gt; b)         -&gt; PushT a         -&gt; PushT b    Append :: Expable b           =&gt; Length           -&gt; PushT b           -&gt; PushT b           -&gt; PushT b    Interleave :: PushT b               -&gt; PushT b               -&gt; PushT b    Reverse :: PushT b            -&gt; PushT b    Rotate :: Length           -&gt; PushT b           -&gt; PushT b </pre>	<pre> apply :: Expable b =&gt; PushT b -&gt; ((Ix -&gt; b -&gt; CM ()) -&gt; CM ()) apply (Generate n ixf) =   \k -&gt; do for_ n \$ \i -&gt;     k i (ixf i)  apply (Use n mem) =   \k -&gt; do for_ n \$ \i -&gt;     k i (cmIndex mem i)  apply (Map f p) =   \k -&gt; apply p (\i a -&gt; k i (f a))  apply (IMap f p) =   \k -&gt; apply p (\i a -&gt; k i (f i a))  apply (Append n p1 p2) =   \k -&gt; apply p1 k &gt;&gt;     apply p2 (\i a -&gt; k (n + i) a)  apply (Interleave p1 p2) =   \k -&gt; apply p1 (\i a -&gt; k (2*i) a) &gt;&gt;     apply p2 (\i a -&gt; k (2*i+1) a)  apply (Reverse p) =   \k -&gt; apply p (\i a -&gt; k ((len p) - 1 - i) a)  apply (Rotate n p) =   \k -&gt; apply p (\i a -&gt; k ((i+n) `mod` (len p)) a) </pre>
--	---

**Figure 6.** The PushT data type and apply function that is obtained from defunctionalization of our Push array API.

## 5. A New Expressive Library

The previous section showed how to defunctionalize Push arrays. But we haven't really gained anything, the library still contains the same functions and they all still do the same thing. Here is the key insight: now that we have a concrete data type instead of a function, we can write new functions on this data type by analysing the value and taking them apart. In particular, we will show how to write functions for our new library which previously belonged in the realm of Pull arrays.

We already know from previous sections that it is possible to convert from Pull arrays to Push arrays, which means that we can also convert from Pull arrays to PushT. If we could also find a way to convert from PushT to Pull arrays it would mean that we can express any function on Pull arrays using PushT. We will demonstrate the conversion by implementing an indexing function.

One characteristic of Push arrays is that in order to look at a specific index, potentially all elements must to be computed. On the defunctionalized Push arrays it is possible to implement an indexing function that is more efficient. No elements other than the one of interest will be computed.

```

index :: Expable a => PushT a -> Ix -> a
index (Generate n ixf) ix = ixf ix
index (Map f p) ix = f (index p ix)
index (Use l mem) ix = cmIndex mem ix
index (IMap f p) ix = f ix (index p ix)
index (Append l p1 p2) ix =
  ifThenElse (ix >* 1)
    (index p2 (ix - 1))
    (index p1 ix)
index (Interleave p1 p2) ix =
  ifThenElse (ix `mod` 2 ==* 0)
    (index p1 (ix `div` 2))
    (index p2 (ix `div` 2))
index (Reverse p) ix =
  index p (len p - 1 - ix)
index (Rotate dist p) ix =
  index p ((ix - dist) `mod` (len p))

```

The implementation of `index` places restrictions on the operations used in the defunctionalization. For example the permutation functions must be invertible. This is another reason for why `ixMap` has been excluded from the language.

Having the `index` function, allows the implementation of a Push to Pull conversion function that does not make the whole array manifest in memory before returning a Pull array.

```

convert :: PushT a -> Pull a
convert p = Pull (\ix -> index p ix) (len p)

```

Being able to index into Push arrays and to convert them to Pull arrays, opens up for implementation of functions that are considered pull also on Push arrays. One such function is `zipWith`.

```
zipWith :: (Expable a, Expable b)
        => (a -> b -> c)
        -> PushT a
        -> PushT b
        -> PushT c
zipWith f a1 a2 =
  generate (min (length a1) (length a2))
    (\i -> f (index a1 i) (index a2 i))
```

Using traditional Push arrays, the `zipWith` function would require one of the input arrays to either be a Pull array or manifest.

## 6. Example Programs and Compilation Output

### 6.1 Fusion

The big benefit of Push arrays is that operations on them fuse automatically. This becomes very clear in the setting of a code generating DSL; just generate the code and count the number of loops. The first example shows that operations are fused. Here an input array is passed through three operations, `map (+1)`, `reverse` and `rotate 3`.

```
ex1 :: (Expable b, Num b) => PushT b -> PushT b
ex1 = rotate 3 . reverse . map (+1)
```

Compiling this program requires that the element type is instantiated, in this case to a `Expr Int`.

```
myVec = CMMem "input" 10

compileEx1 = runCM 0 $
  toVector ((ex1 arr) :: PushT (Expr Int))
  where arr = use myVec
```

The code generated from compiling this program allocates one array, and performs one loop over the input data. This is exactly what we expect from a completely fused program.

```
Allocate "v0" 10 :>>:
For "v1" 10 (
  Write "v0" (((10 - 1) - v1) + 3) % 10) (input[v1] + 1)
)
```

### 6.2 Saxpy

The next compilation example is the `saxpy` operation. This is an operation that we typically would not implement entirely on Push arrays, but this has now been made possible.

```
saxpy :: Expr Float
      -> PushT (Expr Float)
      -> PushT (Expr Float)
      -> PushT (Expr Float)
saxpy a xs ys = zipWith f xs ys
  where
    f x y = a * x + y
```

We compile `saxpy` with two Push arrays that are created by a direct application of `use`.

```
i1 = CMMem "input1" 10
i2 = CMMem "input2" 10

compileSaxpy = runCM 0 $
  toVector (let as = use i1
              bs = use i2
            in saxpy 2 as bs)
```

This results in the following program.

```
Allocate "v0" 10 :>>:
For "v1" 10 (
  Write "v0" v1 ((2.0 * input1[v1]) + input2[v1])
)
```

However, in the case of `saxpy` that uses the `index` function, it is more interesting to see the compiler output when at least one of the arrays is not simply created by a `use`. For example if the first example `ex1` is applied to one of the input arrays, before applying `saxpy`.

```
i1 = CMMem "input1" 10
i2 = CMMem "input2" 10

compileSaxpy = runCM 0 $
  toVector (let as = use i1
              bs = ex1 $ use i2
            in saxpy 2 as bs)
```

In this case the permutations and `map (+1)` from `ex1` is inlined into the indexing into `index2`. This is precisely what the generated code would have looked like if both input arrays had been of Pull array type.

```
Allocate "v0" 10 :>>:
For "v1" 10 (
  Write "v0" v1
    ((2.0 * input1[v1]) +
     (input2[((10 - 1) - ((v1 - 3) % 10))] + 1.0)))
```

## 7. Discussion

This paper answers the question if it is possible to unify Pull- and Push arrays and obtain an array DSL with only a single array type, while maintaining the benefits that Pull- and Push arrays bring separately. By applying defunctionalization to a Push array API we obtain that goal.

The main benefit of creating a concrete data type for Push array programming is that the `index` function can be implemented. This function instantly provides the programmer with all the flexibility of Pull arrays. The crux is that the Push array operations must be safe, permutations need to be invertible and all elements defined.

Here we used defunctionalization as a means to obtain a concrete representation of Push arrays. Now, looking at the API we used as a starting point, coming up with a data type that represents those operations is not hard and could be done in a more ad hoc way. But by using defunctionalization, we cut the amount of thinking necessary to implement compilation of the defunctionalized Push array language down to almost zero.

On a more general note, we have a mantra for dealing with programs written in continuation passing style: “Always defunctionalize the continuation!” Following that mantra almost always yields insights into to program. In some cases a defunctionalized continuation leads to new opportunities to write programs which were not possible before, as we have demonstrated in this paper.

### 7.1 Embedded vs Native

In this paper we have targeted arrays for embedded languages. But what if we wanted to use the new array type natively in a language without any embedding? It is entirely possible to do so but we will have to work harder to provide the kind of fusion guarantees that the embedded language approach provides. The types `Push` and `Pull` are non-recursive and all functions on them are also non-recursive. Achieving fusion for these types are just a matter of inlining and beta-reductions, which are standard optimizations implemented by most compilers. However, the type `PushT` is a recursive type and

the functions manipulating values of this type are also by necessity recursive. In order to achieve fusion for `PushT` we would have to use shortcut fusion or some similar technique [17]. We refrain from going into details here.

## 7.2 Unsafe Index Operations

Some `Push` array libraries [10] contain functions which permutes arrays by transforming indexes, like the following:

```
ixMap :: (Ix -> Ix) -> Push a -> Push a
ixMap f (Push p l) =
  Push (\k -> p (\i a -> k (f i) a)) l
```

These kinds of functions are problematic for our new array library. The problem is that the index transformation function `f`, which permutes the indexes is not guaranteed to be a proper permutation, i.e. a bijection. If we had included `ixMap` in our library we wouldn't have been able to write the `index` function in Section 5, because we would have needed to invert the index transformation function. In our library we have instead opted for a fixed set of combinators which provide specific permutation. This not only solves our problem but we also consider it to be a better library design. The function `ixMap` is a potentially unsafe function and would give undefined results if the programmer were to provide an index transformation function which is not a proper permutation. Another approach to dealing with functions such as `ixMap` would have been to provide a type for permutations which can be inverted and have that as an argument instead of the index transformation function.

## 8. Related Work

### 8.1 Array programming

Many operations can be implemented efficiently on `Pull` arrays and compose without inducing storage of data in memory. This is one of the reasons for why this representation of arrays is being used in many embedded languages. `Feldspar`, is an example of such an embedded language for digital signal processing [6].

Another property of `Pull` arrays is that they are parallelizable. The elements of a `Pull` array are all computed independently and could be computed in any order or in parallel. This property of `Pull` arrays is used in the embedded language `Obsidian`, for general purpose GPU programming [24].

In Pan [15], a similar representation is used for images and in `Repa` [19], the *delayed* array is another example of the same representation. Later versions of `Repa` contains a more refined array representation which allows for efficiently implementing stencil convolutions [? ]. Our line of work is different in that we have chosen to keep the simple `Pull` arrays and add `Push` arrays to be able to efficiently implement stencil computation.

### 8.2 Defunctionalization

Defunctionalization is a technique introduced by Reynolds in his seminal paper on definitional interpreters [22]. The transformation has later been studied by Danvy and Nielsen [12]. Defunctionalization for polymorphic languages has been developed have been developed by two different groups [7, 21], and we make use of these results in this paper.

In a series of papers Olivier Danvy and his co-authors have used defunctionalization and other techniques to establish correspondences between interpreters and abstract machines, among other things [1–4, 8, 9, 11, 14]. In particular, one key step of their correspondence is to defunctionalize continuations to get a first order representation. This is very similar to the work we have presented here, but applied in a different context. However, we go further by looking at the defunctionalized continuation and write new func-

tions on this data type, functions which where not possible to write before.

Another example of defunctionalizing continuations is presented by Fillitre and Pottier [16]. The authors derive a very efficient, first order algorithm for generating all ideals of a forest poset as a Gray code from a functional specification.

## 9. Future Work

`Push` and `Pull` arrays have been central to our research in high performance array programming for a while now. This work furthers our understanding of `Push` arrays, and provides a way to unify functionality of `Pull` and `Push` arrays using a single array representation. However, this implementation of defunctionalized `Push` arrays is just a proof of concept. A natural next step is to replace `Pull`- and `Push` arrays in one of our existing embedded DSLs, `Obsidian` or `Feldspar`, with this single array representation and see how well it fares under those conditions. A natural part of this work would be to generalize the arrays to higher dimensions along the lines of `Repa` [19].

## Acknowledgments

Thanks Michal Palka, Anders Persson, Koen Claessen, Jean-Phillipe Bernardy and Mary Sheeran for important insights, feedback and support.

This research has been funded by the Swedish Foundation for Strategic Research (which funds the Resource Aware Functional Programming (RAW FP) Project) and by the Swedish Research Council.

## References

- [1] M. S. Ager, D. Biernacki, O. Danvy, and J. Midtgaard. A functional correspondence between evaluators and abstract machines. In *Proceedings of the 5th ACM SIGPLAN international conference on Principles and practice of declarative programming*, pages 8–19. ACM, 2003.
- [2] M. S. Ager, D. Biernacki, O. Danvy, and J. Midtgaard. *From interpreter to compiler and virtual machine: a functional derivation*. BRICS, Department of Computer Science, University of Aarhus, 2003.
- [3] M. S. Ager, O. Danvy, and J. Midtgaard. A functional correspondence between call-by-need evaluators and lazy abstract machines. *Information Processing Letters*, 90(5):223–232, 2004.
- [4] M. S. Ager, O. Danvy, and J. Midtgaard. *A functional correspondence between monadic evaluators and abstract machines for languages with computational effects*. BRICS, Department of Computer Science, Univ., 2004.
- [5] J. Ankner and J. D. Svenningsson. An EDSL Approach to High Performance Haskell Programming. In *Proceedings of the 2013 ACM SIGPLAN Symposium on Haskell, Haskell '13*, pages 1–12, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2383-3. . URL <http://doi.acm.org/10.1145/2503778.2503789>.
- [6] E. Axelsson, K. Claessen, M. Sheeran, J. Svenningsson, D. Engdal, and A. Persson. The Design and Implementation of `Feldspar` an Embedded Language for Digital Signal Processing. In *Proceedings of the 22nd international conference on Implementation and application of functional languages, IFL'10*, pages 121–136, Berlin, Heidelberg, 2011. Springer Verlag. ISBN 978-3-642-24275-5. URL <http://dl.acm.org/citation.cfm?id=2050135.2050143>.
- [7] J. M. Bell, F. Bellegarde, and J. Hook. Type-driven defunctionalization. In *Proceedings of the Second ACM SIGPLAN International Conference on Functional Programming, ICFP '97*, pages 25–37, New York, NY, USA, 1997. ACM. ISBN 0-89791-918-1. . URL <http://doi.acm.org/10.1145/258948.258953>.

- [8] M. Biernacka and O. Danvy. A syntactic correspondence between context-sensitive calculi and abstract machines. *Theoretical Computer Science*, 375(1):76–108, 2007.
- [9] D. Biernacki and O. Danvy. *From interpreter to logic engine by defunctionalization*. Springer, 2004.
- [10] K. Claessen, M. Sheeran, and B. J. Svensson. Expressive Array Constructs in an Embedded GPU Kernel Programming Language. In *Proceedings of the 7th workshop on Declarative aspects and applications of multicore programming*, DAMP '12, pages 21–30, New York, NY, USA, 2012. ACM.
- [11] O. Danvy. Defunctionalized interpreters for programming languages. In *Proceedings of the 13th ACM SIGPLAN international conference on Functional programming*, ICFP '08, pages 131–142, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-919-7. . URL <http://doi.acm.org/10.1145/1411204.1411206>.
- [12] O. Danvy and L. R. Nielsen. Defunctionalization at work. In *Proceedings of the 3rd ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, PPDP '01, pages 162–174, New York, NY, USA, 2001. ACM. ISBN 1-58113-388-X. . URL <http://doi.acm.org/10.1145/773184.773202>.
- [13] O. Danvy and L. R. Nielsen. Defunctionalization at work. In *Proceedings of the 3rd ACM SIGPLAN international conference on Principles and practice of declarative programming*, pages 162–174. ACM, 2001.
- [14] O. Danvy, K. Millikin, J. Munk, and I. Zerny. Defunctionalized interpreters for call-by-need evaluation. In *Functional and Logic Programming*, pages 240–256. Springer, 2010.
- [15] C. Elliott, S. Finne, and O. de Moor. Compiling embedded languages. *Journal of Functional Programming*, 13(2), 2003. URL <http://conal.net/papers/jfp-saig/>.
- [16] J.-C. Fillitre and F. Pottier. Producing all ideals of a forest, functionally. *Journal of Functional Programming*, 13(5):945–956, Sept. 2003. URL <http://gallium.inria.fr/~fpottier/publis/filliatre-pottier.ps.gz>.
- [17] T. Harper. A library writer's guide to shortcut fusion. In *Proceedings of the 4th ACM Symposium on Haskell*, Haskell '11, pages 47–58, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0860-1. . URL <http://doi.acm.org/10.1145/2034675.2034682>.
- [18] R. Hughes. A novel representation of lists and its application to the function reverse. *Information processing letters*, 22(3):141–144, 1986.
- [19] G. Keller, M. M. Chakravarty, R. Leshchinskiy, S. Peyton Jones, and B. Lippmeier. Regular, shape-polymorphic, parallel arrays in haskell. In *Proceedings of the 15th ACM SIGPLAN international conference on Functional programming*, ICFP '10, pages 261–272, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-794-3. . URL <http://doi.acm.org/10.1145/1863543.1863582>.
- [20] A. Kulkarni and R. R. Newton. Embrace, Defend, Extend: A Methodology for Embedding Preexisting DSLs, 2013. Functional Programming Concepts in Domain-Specific Languages (FPCDSL'13).
- [21] F. Pottier and N. Gauthier. Polymorphic typed defunctionalization. In *Proceedings of the 31st ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, POPL '04, pages 89–98, New York, NY, USA, 2004. ACM. ISBN 1-58113-729-X. . URL <http://doi.acm.org/10.1145/964001.964009>.
- [22] J. C. Reynolds. Definitional interpreters for higher-order programming languages. In *Proceedings of the ACM annual conference - Volume 2*, ACM '72, pages 717–740, New York, NY, USA, 1972. ACM. . URL <http://doi.acm.org/10.1145/800194.805852>.
- [23] J. Svenningsson and E. Axelsson. Combining Deep and Shallow Embedding for EDSL. In H.-W. Loidl and R. Pea, editors, *Trends in Functional Programming*, volume 7829 of *Lecture Notes in Computer Science*, pages 21–36. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-40446-7. . URL [http://dx.doi.org/10.1007/978-3-642-40447-4\\_2](http://dx.doi.org/10.1007/978-3-642-40447-4_2).
- [24] J. Svensson, M. Sheeran, and K. Claessen. Obsidian: A Domain Specific Embedded Language for Parallel Programming of Graphics Processors. In S.-B. Scholz and O. Chitil, editors, *Implementation and*

*Application of Functional Languages*, volume 5836 of *Lecture Notes in Computer Science*, pages 156–173. Springer Berlin Heidelberg, 2011. ISBN 978-3-642-24451-3. .